

CRISNER: A Practically Efficient Reasoner for Qualitative Preferences

Ganesh Ram Santhanam¹, Samik Basu¹, and Vasant Honavar²

¹Iowa State University, Ames, Iowa 50011, USA

²Pennsylvania State University, University Park, PA 16801, USA
 {gsanthan,sbasu}@iastate.edu; vhonavar@ist.psu.edu

Abstract. We present CRISNER (Conditional & Relative Importance Statement Network PrEference Reasoner), a tool that provides practically efficient as well as exact reasoning about qualitative preferences in popular ceteris paribus preference languages such as CP-nets, TCP-nets, CP-theories, etc. The tool uses a model checking engine to translate preference specifications and queries into appropriate Kripke models and verifiable properties over them respectively. The distinguishing features of the tool are: (1) exact and provably correct query answering for testing dominance, consistency with respect to a preference specification, and testing equivalence and subsumption of two sets of preferences; (2) automatic generation of proofs evidencing the correctness of answer produced by CRISNER to any of the above queries; (3) XML inputs and outputs that make it portable and pluggable into other applications. We also describe the extensible architecture of CRISNER, which can be extended to new preference formalisms based on ceteris paribus semantics that may be developed in the future.

Keywords: Qualitative Preferences, Tool, CP-net, Model Checking

1 Introduction

Several qualitative preference reasoning languages have been developed in the last two decades, such as CP-nets, TCP-nets, CP-Theories, etc. Despite their relatively high expressive power (compared to quantitative preference formalisms), their widespread application and use in practice is limited, at least in part due to their hardness. The basic reasoning tasks such as dominance and consistency testing are known to be PSPACE-complete for a simple language such as the CP-net. Another reasoning task, checking whether the equivalence or subsumption of the preferences induced by one agent with respect to those induced by another, is important in multi-agent scenarios and negotiation but also known to be hard [17]. Past works to cope with this hardness include restricting the expressivity of the languages to obtain tractable fragments, and heuristics that yield results in acceptable time but not guaranteed to be exactly correct. Nevertheless, there are applications such as negotiation, planning, security policies, etc. that call for exact reasoning about qualitative preferences with guarantees of correctness (e.g., choosing the best policy to defend a computer network).

In this paper, we present CRISNER (Conditional & Relative Importance Statement Network PrEference Reasoner) [14], a tool that provides practically efficient as well as exact reasoning about qualitative preferences in popular *ceteris paribus*¹ [10] preference languages such as CP-nets, TCP-nets and CP-theories. For a preference specification P consisting of a set $\{p_1, p_2, \dots, p_n\}$ of qualitative preference statements in any of the above CP-languages², the *ceteris paribus* semantics for dominance, consistency etc. are given in terms of reachability over an *induced preference graph* wherein each node corresponds to an outcome and each edge from one node to another represents a preference from the latter node to the former node induced by some statement p_i in P . CRISNER uses a model checking [8] engine NuSMV [6] to translate preference specifications and queries into appropriate Kripke structure models [11] and reachability properties over them respectively.

Encoding Preferences as Kripke models. Given a specification P , CRISNER first succinctly encodes the induced preference graph (IPG) of P into a Kripke structure model K_P in the language of the NuSMV model checker. Although Kripke structures are typically used to represent semantics of temporal and modal logics, we leverage earlier work [15,12] that demonstrated their use for encoding preference semantics. For testing dominance and consistency with respect to each P , CRISNER generates the model K_P only once. Subsequently for each preference query q posed against P , CRISNER translates q into a temporal logic formula φ_q in computation-tree temporal logic (CTL) [13,7] such that $K_P \models \varphi_q$ if and only if q holds true according to the *ceteris paribus* semantics of P . CRISNER then queries the model checker with the model K_P and φ_q which either affirms q or returns false with a counterexample. For answering queries related to equivalence and subsumption checking of two sets of preferences P_1 and P_2 , CRISNER constructs a combined IPG $K_{P_{12}}$ and uses temporal queries in CTL to identify whether every dominance that holds in P_1 also holds in P_2 and vice-versa [17].

Justification of Query Answers. The answers to queries computed by CRISNER are exact and provably correct for dominance, consistency, equivalence and subsumption queries. Because CRISNER uses the model checker for answering queries, CRISNER is also able to provide proofs or justifications to queries that returned **false**. CRISNER automatically builds proofs evidencing why the query did not hold true, by collecting and examining the model checker’s counterexample and producing a sequence of preference statements whose application proves the correctness of CRISNER’s result.

Tool Architecture. CRISNER is developed in pure Java and is *domain agnostic* in the sense that any set of variables with any domain can be included in a preference specification, although it is optimized for variables with binary domains. It accepts preference specifications and queries in a XML format, which provides a common and generic syntax using which users can specify preferences

¹ *Ceteris paribus* is a Latin term for “all else being equal”.

² Henceforth, we will refer to the languages CP-nets, TCP-nets and CP-theories collectively as ‘CP-languages’ for brevity.

for CP-languages. The results (answers and proofs) for the corresponding queries are also saved in the form XML, so that the results can be further transformed into vocabulary that is more easily understandable by domain users. We describe the architecture of CRISNER and how it can be extended to other *ceteris paribus* preference formalisms that may be developed in the future.

CRISNER has been in development for over two years, and to our knowledge, CRISNER is one of the first attempts to develop practical tools for hard qualitative preference reasoning problems. We hope that CRISNER inspires the use of qualitative preference formalisms in practical real world applications, and the development of further qualitative preference languages.

2 Background: Syntax and Semantics of CP-languages

Let $X = \{x_i \mid 0 < i \leq n\}$ be a set of preference variables or attributes. For each $x_i \in X$ let D_i be the set of possible values (i.e., domain) such that $x_i = v_i \in D_i$ is a valid assignment to the variable x_i . We use Φ, Ω (indexed, subscripted or superscripted as necessary) to denote subsets of X . The set $\mathcal{O} = \{\alpha \mid \prod_{x_i \in X} D_i\}$ of assignments to variables in X is the set of *alternatives*. The *valuation* of an alternative $\alpha \in \mathcal{O}$ with respect to a variable $x_i \in X$ is denoted by $\alpha(x_i) \in D_i$.

2.1 Preference Relations, Statements & Specifications

Given a set \mathcal{O} of n alternatives, a direct specification of a binary preference relation \succ over \mathcal{O} is difficult, as it requires the user to compare up to $O(n^2)$ pairs of alternatives, which is prohibitive in time. Hence, many preference languages allow for succinct specification of the preference relation over alternatives in terms of *preference relations over the set of attributes that describe the alternatives and their respective valuations or domains*.

Preference Relations Qualitative preference relations can be either (a) *intra-variable preference relations over valuations* of an attribute; or (b) *relative importance preference relations over attributes*. For any $\Phi \subseteq X$, we will use the notation \succ_Φ to denote a preference relation over D_Φ , the set of partial assignments to attributes in Φ . For a single attribute $x_i \in X$, the intra-attribute preference relation over its valuations (D_i) will be denoted by $\succ_{\{x_i\}}$ or alternatively \succ_i . For example, to formally specify that the valuation v_i is preferred to the valuation v'_i for attribute x_i where $v_i, v'_i \in D_i$, we will write $x_i = v_i \succ_i x_i = v'_i$. We will use the notation \triangleright to indicate relative importance between attributes or between sets of attributes.

Preference Specifications & Statements In the CP-languages, preferences are expressed in terms of a preference specification, which is a set $P = \{p_i\}$ of preference statements. Each statement p may specify a binary relation over the set X of preference variables (relative importance) and/or a binary relation over the domain of a particular variable (intra-variable preference). The syntax of the preference statements is given below.

CP-nets [2,9] allow the specification of only conditional intra-variable preferences³; TCP-nets [5] allow the users to specify *pairwise* relative importance

³ We use the term CP-nets to refer to the more general formalism defined by Goldsmith et al. [9].

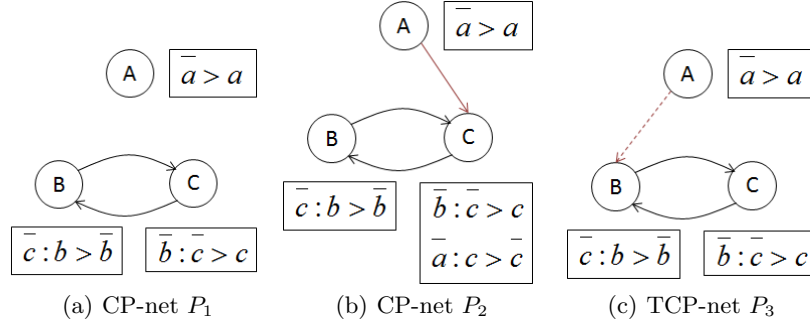


Fig. 1. Consistent and inconsistent CP-nets; TCP-net

among variables in addition to conditional intra-variable preferences as in CP-nets. CP-Theories [19] extend TCP-nets by further allowing the specification of the relative importance of one variable over a *set* of variables conditioned on another set of variables. As CP-theories strictly generalize CP-nets and TCP-nets [19], we give the syntax for CP-theories here. A CP-Theory consists of statements of the form

$$\varrho : x_i = v_i \succ_i x_i = v'_i [\Omega]$$

where ϱ is an assignment to the set $\Phi \subseteq X$ of variables that defines the condition under which the preference holds, $v_i, v'_i \in D_i$, $\Omega \subseteq X$, and $\Phi, \Omega, \{x_i\}$ and $(X - \Phi - \Omega - \{x_i\})$ are disjoint. The statement expresses the relative importance of the variable x_i over the set Ω of variables under the condition ϱ . Note that CP-nets can be expressed as CP-Theories by fixing $\Omega = \emptyset$ (i.e., $|\Omega| = 0$); and TCP-nets can be expressed as CP-Theories by fixing $|\Omega| = 0$ or 1. Hence we use the above general syntactic form to refer to a preference statement in any of the CP-languages. Figure 1 shows two CP-nets P_1 and P_2 , and a TCP-net P_3 where the red dotted arrow from A to B indicates that A is more important than B .

2.2 Ceteris Paribus Semantics

The ceteris paribus semantics of the CP-languages define an induced preference graph with nodes corresponding to the outcomes, and the edges induced by the language-specific interpretation of preference statements. According to the *ceteris paribus* interpretation [10,3], each preference statement $p \in P$ allows a set of changes to the valuation(s) of one or more variables in an alternative β in order to obtain a more preferred alternative α , while other variables remain fixed. Such a change is called an **improving flip**. The improving flips induced by the different types of preference statements is summarized below; we do not elaborate on the semantics of each CP-language, and refer the interested reader to [2,5,18] for details of the semantics of the respective languages.

Consider an intra-variable preference statement p of the form $\varrho : v_i \succ_{\{x_i\}} v'_i$, which can be specified in all the languages we consider. Given two alternatives $\alpha, \beta \in \mathcal{O}$, the ceteris paribus interpretation of the statement p induces an improving flip $(\beta, \alpha) \in E$ in $\delta(P)$ if α and β differ only in x_i , and $\alpha(x_i) = v_i$ and

$\beta(x_i) = v'_i$. In other words for any statement p , the valuation of only one variable can be flipped at a time, and that to a more preferred valuation with respect to p , while other variables remain fixed.

Definition 1 (Preference Semantics for Intra-attribute Preference [2]). *Given an intra-attribute preference statement p in a preference specification P of the form $\varrho : x_i = v_i \succ_{\{x_i\}} x_i = v'_i$ where $\varrho \in D_\Phi, \Phi = \rho(x_i) \subseteq X$ and two alternatives $\alpha, \beta \in \mathcal{O}$, there is an improving flip from β to α in $\delta(P)$ induced by p if and only if*

1. $\exists x_i \in X : \alpha(x_i) = v_i$ and $\beta(x_i) = v'_i$,
2. $\forall x_j \in \Phi : \alpha(x_j) = \beta(x_j) = \varrho(x_j)$, and
3. $\forall x_k \in X \setminus \{x_i\} \setminus \Phi : \alpha(x_k) = \beta(x_k)$.

In the above definition, the first condition arises from the intra-attribute preference statement x_i ; the second condition enforces the condition ϱ in p that states that α and β should concur on the parent variables of x_i ; and the third enforces the ceteris paribus condition that states that α and β should concur on all the other variables.

TCP-nets and CP-Theories allow the specification of relative importance preference of one variable over one or more variable respectively. Hence, multiple variables can change in the same improving flip because a statement of relative importance of one attribute over others means that the user is willing to improve the valuation of the more important attribute *at the expense* of worsening the less important attribute(s).

Definition 2 (Preference Semantics for Relative Importance of one Attribute over a Set [5,18]). *Given a relative importance preference statement p in a preference specification P of the form $\varrho : x_i = v_i \succ_{\{x_i\}} x_i = v'_i [\Omega]$ where $\varrho \in D_\Phi, \Phi \subseteq X$ and two alternatives $\alpha, \beta \in \mathcal{O}$, there is an improving flip from β to α in $\delta(P)$ induced by p if and only if*

1. $\exists x_i \in X : \alpha(x_i) = v_i$ and $\beta(x_i) = v'_i$,
2. $\forall x_j \in \Phi : \alpha(x_j) = \beta(x_j) = \varrho(x_j)$, and
3. $\forall x_k \in X \setminus \{x_i\} \setminus \Omega \setminus \Phi : \alpha(x_k) = \beta(x_k)$.

In the above definition, the first condition arises from the preference statement p on x_i ; the second condition enforces the condition ϱ in p ; and the third enforces the ceteris paribus condition and allows for unrestricted changes to the attributes that are less important than x_i (in Ω) when this preference statement is applied.

The above interpretations define valid improving flips induced by statements in P that correspond to edges in the **induced preference graph** denoted $IPG(P)$, whose nodes are the set of all outcomes or alternatives (i.e., the set of all assignments to all preference variables), and which represents the dominance relation \succ_P over the outcomes. Figure 2 shows the induced preferences graphs for the respective CP-nets and TCP-net in Figure 1. Note that the red solid edges in $IPG(P_2)$ are those induced by the dependency of C on A in P_2 . Similarly, the red dotted edges in $IPG(P_3)$ are those induced by the relative importance of A over B in P_3 .

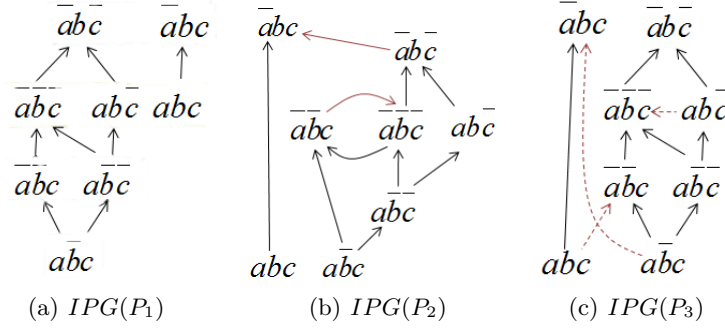


Fig. 2. Induced preference graphs of CP-nets and TCP-net

3 Preference Queries

Computing answers to preference queries with respect to a given preference specification in the ceteris paribus semantics amounts to making querying properties related to reachability on the induced preference graph. We consider the following preference queries that have been implemented in CRISNER in this paper.

Definition 3 (Dominance & Consistency Testing [2,9,19]). *Given a preference specification P consisting of a set of preference statements $\{p_1, p_2 \dots p_n\}$, and two outcomes $\alpha, \beta \in \mathcal{O}$*

1. **Dominance Testing** ($\alpha \succ_P \beta$) *asks whether there a sequence of improving flips from β to α in $IPG(P)$?*
2. **Consistency Testing** *asks whether the preferences induced by P are consistent, i.e., is there is a cycle in $IPG(P)$?*

Definition 4 (Preference Equivalence & Subsumption [17]). *Given two preference specifications P_1, P_2 and two outcomes $\alpha, \beta \in \mathcal{O}$,*

1. **Preference Subsumption** ($P_1 \sqsubseteq P_2$) *asks whether $\alpha \succ_{P_1} \beta \Rightarrow \alpha \succ_{P_2} \beta$.*
2. **Preference Equivalence** ($P_1 \equiv P_2$) *asks whether $P_1 \sqsubseteq P_2$ and $P_2 \sqsubseteq P_1$.*

3.1 Preference Query Answering via Model Checking

In order to address the PSPACE-hardness of the problems of dominance and consistency, CRISNER implements the model checking based approach to preference query reasoning as presented in the series of works by Santhanam et al. [15,16,17]. There are two direct benefits of using the model checking approach to answering preference queries. First, by using the model encoding techniques presented in the above works, then dominance and consistency queries can be transformed into equivalent reachability queries on the Kripke model in a straightforward way. The transformation from the preference specification and query to the Kripke model and temporal logic formula respectively is at a syntactic level prior to execution of the model checker (without having to build the induced preference graph); thus preserving the benefits of succinctness of the CP-languages.

When the query is actually executed on the model checker, the induced preference graph is expanded to the extent needed by the model checker to verify the temporal logic formula. This would enable us to leverage the decades of advances in model checking algorithms and tools. The second direct benefit if the model checker returns a counter example to the encoded temporal logic formula (corresponding to the original preference query), which is in terms of states and transition sequences in the Kripke model, then it can be conveniently mapped back to nodes and paths in the induced preference graph. This allows CRISNER to automatically generate justifications for dominance and consistency queries by producing sequences of improving flips in the induced preference graph that either prove the dominance or disprove the consistency.

CTL model checking We use formulas in computation-tree temporal logic (CTL) [8] for verifying reachability within the Kripke model generated by CRISNER. Our choice of model checker is NuSMV [6], an open source and widely used tool. CTL is an extension of propositional logic; CTL uses propositional and temporal connectives to express temporal properties, whose semantics is given in terms of a set of states in the Kripke structure where the properties are satisfied. We briefly outline the syntax and semantics of some CTL connectives below, and refer the reader to [8] for details. The syntax of CTL is described as follows:

$$\varphi \rightarrow true \mid \text{Atomic-Propositions} \mid \neg\varphi \mid \varphi \vee \varphi \mid EX\varphi \mid EG\varphi \mid E(\varphi \cup \varphi)$$

The semantics of a CTL formula φ is given in terms of the set of states in a Kripke structure that satisfy φ . The propositional constant *true* is satisfied in all states; the proposition p holds in states labeled with p . The negation of φ is satisfied if the formula does not hold. The disjunct of two CTL formulas is satisfied by states if at least one of the disjuncts is satisfied. The rest of the operators in the CTL syntax are temporal operators that quantify the states and the transitions. The property $EX\varphi$ is satisfied in any state s which can reach some (E, for existential quantification) state t in one transition (X, for one step reachability) such that t satisfies φ . The property $EG\varphi$ is satisfied in any state s which has some (E) path where every (G) state t_i 's in the path satisfy φ . Finally, the property $E(\varphi_1 \cup \varphi_2)$ is satisfied in any state s which has some (E) path where there exists a state t which satisfied φ_2 and in all states before t , φ_1 is satisfied. Semantics for other CTL formulas are described in terms of the above, e.g., $EF\varphi$ is satisfied in any state from where there exists a path eventually leading to a state that satisfies φ . This is equivalent to $E(true \cup \varphi)$. There are other temporal connectives which we mention and explain as needed in terms of the above in this paper.

3.2 Encoding Preference Queries in CTL

Here we outline the overall strategy to obtain answers to preference queries in the languages CP-languages. We assume that CRISNER is given a preference specification P and a preference query q that is a test for either a dominance, consistency, preference subsumption or preference equivalence. The task of CRISNER

is then to compute whether q holds (or not) with respect to the ceteris paribus semantics of P .

Dominance Testing Testing dominance of an outcome over another amounts to simply checking reachability from one outcome to another in $IPG(P)$. To answer the dominance query $\alpha \succ_P \beta$, CRISNER constructs a temporal logic formula that corresponds to reachability from β to α in $IPG(P)$ and executes the query on the NuSMV [6] model checker, affirming dominance if and only if the model checker returns true. Defining AP to be the set of preference variables in P and φ_α to be the formula encoding the set of variables assigned *true* in the outcome α , the above dominance can be encoded in the following CTL formula:

$$\varphi_\beta \Rightarrow \mathbf{EF} \varphi_\alpha$$

Consistency Testing Consistency testing amounts to checking that there are no cycles in $IPG(P)$. To test consistency for P , CRISNER verifies the following CTL formula, which states that for all **start** states (If no start states are initialized, NuSMV non-deterministically considers all states in the Kripke model as start states), there must be no reachable node from which there is a path back to α , i.e. there is no cycle in $IPG(P)$.

$$\mathbf{start} \Rightarrow \neg \mathbf{EX} (g = 1 \wedge \mathbf{EF} \mathbf{start})$$

In the above, $g = 1$ holds whenever the destination state of a current state results from an improving flip as per the underlying preference semantics.

Preference Subsumption and Equivalence Preference equivalence (subsumption) checking [17] amounts to testing whether the preferences induced by two preference specifications P_1 and P_2 are the same (or such that one subsumes the other) or not. Given P_1 and P_2 , CRISNER constructs an SMV model $K_{P_{12}}$ of the *combined induced preference graph* (CIPG) [17] that encodes the union of preferences induced by P_1 's preference statements and the reverse of those induced by P_2 's preference statements. In order to answer whether P_1 subsumes the preferences induced by P_2 or not, CRISNER constructs the following temporal logic formula that is verified by NuSMV if and only if P_1 subsumes P_2 (we use g_1 and g_2 to refer to the g for the models corresponding to P_1 and P_2 respectively).

$$\varphi : \mathbf{AX} (g_1 \Rightarrow \mathbf{EX} \mathbf{E} [g_2 \mathbf{U} (\mathbf{start} \wedge g_2)])$$

The above states that whenever there is an improving flip in $IPG(P_1)$ from α to an outcome $(\mathbf{AX} g_1..)$, then there exists a sequence of improving flips $(\mathbf{EXE}(g_2 \mathbf{U}...))$ from that outcome in $IPG(P_2)$ back to α . In the above, all possible α is captured by the proposition **start**. Preference equivalence is checked by in turn verifying that $P_1 \sqsubseteq P_2$ and $P_2 \sqsubseteq P_1$.

4 XML Input Language

CRISNER accepts a preference specification for any of the CP-languages in an XML format. The preference specification consists of a declaration of the

preference variables, their domains and a set of preference statements. Each preference statement is of the form discussed in Section 2.1, and expresses an intra-variable and/or relative importance preference relation over the domain of a variable.

4.1 Defining Preference Variables

Figure 3 shows part of a preference specification defining variables and their domains. The preference variable a has a binary domain with values 0 and 1, whereas x has a domain $\{0, 1, 2\}$. Note that CRISNER supports domain valuations with string values that are combinations of letters and numbers, as allowed by NuSMV.

<pre style="margin: 0;"><VARIABLE> <NAME>a</NAME> <DOMAIN-VALUE>0</DOMAIN-VALUE> <DOMAIN-VALUE>1</DOMAIN-VALUE> </VARIABLE></pre>	<pre style="margin: 0;"><VARIABLE> <NAME>x</NAME> <DOMAIN-VALUE>0</DOMAIN-VALUE> <DOMAIN-VALUE>1</DOMAIN-VALUE> <DOMAIN-VALUE>2</DOMAIN-VALUE> </VARIABLE></pre>
(a)	(b)

Fig. 3. XML encoding of definitions of preference variable a with domain size 2 and 3

4.2 Specifying Conditional Preference Statements

The listing in Figure 4 shows a portion of a preference specification that declares preferences over values of the variable c conditioned on the variables b and a respectively. The **VARIABLE** tag identifies the variable whose preferences are being specified. Note that there can be multiple conditions or no conditions as well. In addition, there can also be multiple preferences for a variable, e.g., if there is a variable with domain of 0, 1, 2 then to specify the total order $0 \succ 1 \succ 2$ one would encode $0 \succ 1$ as one preference followed by $1 \succ 2$. CRISNER requires that the variable names and their assignments match with the preference variable declarations in the file; otherwise the tool reports an appropriate error stating that the variable is not defined in the preference specification.

4.3 Specifying Relative Importance Preferences

In order to allow specification of relative importance of one variable over another, as in a TCP-net, CRISNER allows the tag **REGARDLESS-OF** within a preference statement. Figure 5(a) declares a preference statement that says (in addition to the fact that $a = 0 \succ_a a = 1$) that a is relatively more important than b . In order to specify relative importance of one variable over a set of other variables (simultaneously) as allowed by a CP-theory, the user can specify multiple **REGARDLESS-OF** tags within the same preference statement. For instance, Figure 5(b) shows a preference statement that declares that a is relatively more important than $\{b, c\}$.

5 Encoding Preferences as SMV Models

CRISNER encodes Kripke models for a preference specification as described earlier. Now we discuss constructs in NuSMV used by CRISNER for the encoding.

<pre> <PREFERENCE-STATEMENT> <STATEMENT-ID>p3</STATEMENT-ID> <VARIABLE>c</VARIABLE> <CONDITION>b=0</CONDITION> <PREFERENCE>0:1</PREFERENCE> </PREFERENCE-STATEMENT> </pre> <p style="text-align: center;">(a)</p>	<pre> <PREFERENCE-STATEMENT> <STATEMENT-ID>p4</STATEMENT-ID> <VARIABLE>c</VARIABLE> <CONDITION>a=0</CONDITION> <PREFERENCE>1:0</PREFERENCE> </PREFERENCE-STATEMENT> </pre> <p style="text-align: center;">(b)</p>
---	---

Fig. 4. XML encoding of conditional preference statements p_3 and p_4 in a CP-net

<pre> <PREFERENCE-STATEMENT> <STATEMENT-ID>p1</STATEMENT-ID> <VARIABLE>a</VARIABLE> <PREFERENCE>0:1</PREFERENCE> <REGARDLESS-OF>b</REGARDLESS-OF> </PREFERENCE-STATEMENT> </pre> <p style="text-align: center;">(a)</p>	<pre> <PREFERENCE-STATEMENT> <STATEMENT-ID>p1</STATEMENT-ID> <VARIABLE>a</VARIABLE> <PREFERENCE>0:1</PREFERENCE> <REGARDLESS-OF>b</REGARDLESS-OF> <REGARDLESS-OF>c</REGARDLESS-OF> </PREFERENCE-STATEMENT> </pre> <p style="text-align: center;">(b)</p>
---	--

Fig. 5. XML encoding of conditional preference statements p_3 and p_4 in a CP-net

5.1 Encoding Preference Variables & Auxiliary Variables

In order to encode the CP-net P_1 in our earlier example, CRISNER generates the code for the SMV model as shown in Figure 6. We explain the translation of a preference specification into an SMV model by CRISNER through this example.

CRISNER defines just the `main` module, with 3 variables corresponding to the preference variables in P_1 and another variable g , which we will explain shortly. We overload a, b, c to refer to variables in the Kripke model and variables in the preference specification, hence valuations of a, b, c in a state s in the Kripke model respectively correspond to the valuations of the preference variables a, b, c in the preference specification P . The variables a, b, c are *state* variables in the SMV model i.e., their valuations stored by the model checker for each state explored during model checking.

The `IVAR` variables cha, chb, chc are modeled as *input* variables, i.e., their valuations are not stored as part of each state. The model checker initializes them non-deterministically for each state so that all paths are open for exploration by the model checker during verification. Each preference statement is translated into an appropriate guard condition for a transition in the Kripke model, and the semantics of variables cha, chb, chc either allows or disallows the change in the value of the corresponding preference variable a, b, c , in accordance with the improving flip semantics.

Identifying transitions corresponding to improving flips. The additional g variable is defined to be *true* exactly when the model checker transitions from a state corresponding to one outcome to a state corresponding to another outcome (not transitions between states corresponding to the same outcome). Hence, we can conveniently refer to transitions in the Kripke models that correspond to improving flips by constraining g to have valuation 1 in the destination state.

Referencing start states explored by NuSMV. The FROZEN variables a_0 , b_0 , c_0 are constrained to be fixed with the values of the variables a, b, c respectively at the start of the model checking algorithm via the DEFINE and INIT constructs. This allows us to refer to the state non-deterministically picked by the model checker as the start of model exploration using *start*. This is used for computing consistency, preference subsumption and preference equivalence.

<pre> MODULE main VAR a : {0,1}; c : {0,1}; b : {0,1}; g : {0,1}; FROZENVAR a_0 : {0,1}; b_0 : {0,1}; c_0 : {0,1}; IVAR cha : {0,1}; chb : {0,1}; chc : {0,1}; DEFINE start := a=a_0 & b=b_0 & c=c_0; INIT start=TRUE; </pre>	<pre> ASSIGN next(a) := case a=1 & cha=1 & chb=0 & chc=0 : 0; TRUE : a; esac; next(c) := case c=1 & b=0 & cha=0 & chb=0 & chc=1 : 0; TRUE : c; esac; next(b) := case b=0 & c=0 & cha=0 & chb=1 & chc=0 : 1; TRUE : b; esac; next(g) := case a=1 & cha=1 & chb=0 & chc=0 : 1; c=1 & b=0 & cha=0 & chb=0 & chc=1 : 1; b=0 & c=0 & cha=0 & chb=1 & chc=0 : 1; TRUE : 0; esac; </pre>
---	---

Fig. 6. SMV code for Kripke Model encoding $IPG(P_1)$

5.2 Encoding Preference Statements

Encoding Intra-variable Preferences. To encode a intra-variable preference statement for a variable x with a condition ρ on the other variables, the `next(x)` construct encodes guards such that the valuations of the other variables correspond to those in the condition ρ , and valuation of chx is 1 while all other ch variables are set to 0. As an example, `next(a)` includes a transition such that c changes from 1 to 0 precisely when $b = 0$ and $chc = 1$ ($cha = 0, chb = 0$, allowing only c to change in that transition), which corresponds to the improving flip induced by p_3 conditioned on the value of a in the CP-net P_1 (Figure 6).

Encoding Relative Importance. For modeling relative importance preference statements, multiple IVAR variables can be assigned 1 in guard conditions such that the more important and less important preference variables can change in the same transition - corresponding to an improving flip for relative importance. For example, Figure 7 shows a snippet of the SMV code that models the transitions arising from the relative importance of a over b as in the TCP-net P_3 shown in Figure 1(c). Note that cha and chb are set to 1 for the second guard condition of `next(b)`, allowing a to change to a preferred value trading off b . In order to model relative importance as in a CP-theory where one variable is more important than multiple others, a similar encoding is used, except that all the corresponding ch variables are set to 1.

```

next(b) := case
  b=0 & c=0 & cha=0 & chb=1 & chc=0 : 1;
  a=1 & cha=1 & chb=1 & chc=0 : {0,1};
  TRUE : b;
esac;

```

Fig. 7. Encoding relative importance preferences for the TCP-net P_3

5.3 Justification of Query Results

In addition to answering preference queries posed against preference specifications, CRISNER also provides a justification of the result where appropriate. In order to obtain justification, CRISNER uses the counterexamples returned by NuSMV model checker whenever a temporal logic formula is not satisfied.

Extracting a Proof of Dominance. In the case of a dominance query, if CRISNER returns **true**, we construct another temporal logic formula that states the negation of the dominance relationship, which obtains a sequence of outcomes corresponding to an improving flipping sequence from the lesser preferred to the more preferred outcome from the model checker. Suppose that we want to obtain proof that an alternative α dominates another alternative β . This means that $\varphi_\beta \rightarrow \text{EF}\varphi_\alpha$ holds. We then verify $\neg(\varphi_\beta \rightarrow \text{EF}\varphi_\alpha)$, which obtains a sequence of states in the Kripke model corresponding to an improving flipping sequence from β to α from the model checker corresponding to an improving flipping sequence from β to α which serves as the proof of dominance.

Extracting a Proof of Inconsistency. In the case of a consistency query (see Section 3.2), CRISNER returns a sequence of outcomes corresponding to an improving flipping sequence from an outcome to itself (indicating a cycle in the induced preference graph) whenever the preference specification input is inconsistent.

Extracting a Proof of Non-subsumption. For a preference subsumption query $P_1 \sqsubseteq P_2$, CRISNER provides an improving flip from one outcome to another induced by P_1 but not induced by P_2 whenever the query does not hold.

In the above, counterexamples returned by NuSMV are in terms of states and transitions in the Kripke model; CRISNER parses and transforms the counterexamples back into a form that relates to the preference variables, outcomes and improving flips in the induced preference graph of the preference specification, and saves it in an XML format.

6 Architecture

CRISNER is built using the Java programming language⁴. The architecture of CRISNER consists of several components as depicted in Figure 6.

The XML parser is used to parse the preference specifications and preference queries input by the user⁵. The CP-language translator is a critical component

⁴ Third party libraries used by CRISNER are listed in the project site [14].

⁵ While currently CRISNER does not use XML schema or DTD to validate the XML input, we plan to enforce that in future.

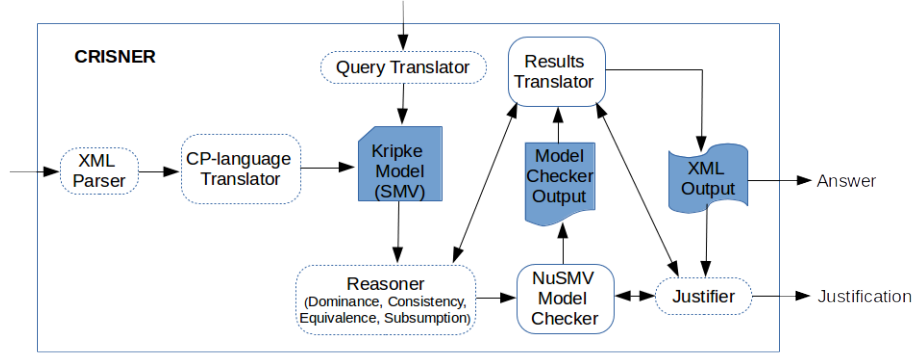


Fig. 8. Architecture and components of CRISNER preference reasoner

that constructs the SMV code for the Kripke model for the preference specification input. It declares the necessary variables with their domains, sets up the **DEFINE**, **TRANS** and **INIT** constraints and finally generates guard conditions for enabling transitions corresponding to improving flips induced by each preference statement (as discussed in Section 3.1).

CRISNER provides two interfaces for preference reasoning. The first is a simple command line menu-drive console interface where the user can provide either (a) one preference specification as input and then use the menu to pose dominance and consistency queries, or (b) two preference specifications as input and pose a subsumption or equivalence query. The answer (**true/false**) obtained and the justification for the answer (where possible) is provided on the console. Another way of using CRISNER is to specify preference queries in an XML file that contains a dominance or consistency or preference equivalence or subsumption query, and identifies the preference specification against which the query should be executed. The query translator component parses queries specified in XML format and provides it to the Reasoner component. Further details about the XML tags and examples of preference specification, preference queries and sample SMV code generated by CRISNER are available from CRISNER's project website <http://www.ece.iastate.edu/~gsanthan/crisner.html>.

The Reasoner is another critical component in CRISNER that constructs a temporal logic formula corresponding to the preference query posed by the user, and invokes the NuSMV model checker to verify the formula. The result and any counter examples generated by the model checker are parsed by the Results Translator, and saved in XML format by the XML Encoder. If a counter example is applicable to the preference query, then the Justifier parses the XML output and executes any followup queries on the model checker (e.g., verification of the negation of the dominance query) to provide the user with the appropriate proof.

Extensibility. Although CRISNER currently supports the CP-net, TCP-net and CP-theory formalisms, it can be extended to support another qualitative preference language, as long as the semantics of the language is described in terms of an induced preference graph. To extend support, the XML parser must be extended to support the syntax of the new language; and the CP-language translator must be extended to generate SMV code according to the semantics of

the new language. We are currently working on including support for the conditional importance network (CI-net) [4] language for representing and reasoning with preferences in the context of fair division of indivisible goods. As another example, one can think of a new preference formalism may allow expression of preferences of one partial assignment of variables over another. In such a case, the CP-translator component can be extended to include a translation for such a preference statement into guard conditions in the Kripke model representing the induced preference graph. CRISNER can also be extended to support new preference queries, for example, computing a weak order or total order extension of the partial order induced by a preference specification. In this case the Query translator and Reasoner should be extended to translate the new queries into one or more appropriate temporal logic formulas (as described in [12]) and the Justifier should be extended to construct and execute follow up queries that obtain proofs for the answers.

Scalability. While we have not yet performed a systematic experiment studying CRISNER’s runtime performance for preference specifications of different sizes (number of preference variables), our preliminary tests have revealed that CRISNER answers dominance (including the computation of justification for a when applicable) in less than a minute on average for upto 30 variables on a 8GB Corei7 Windows 7 desktop. Although CRISNER allows variables with domain size $n > 2$, the model checking performance degrades quickly with increasing n ; this can be alleviated by configuring the NuSMV model checker to use multi-way decision diagrams [1].

7 Concluding Remarks

We presented CRISNER, a tool for specifying and reasoning with qualitative preference languages such as CP-net, TCP-net and CP-theory. CRISNER translates preference specifications and queries with respect to them provided in XML format into a Kripke structure and corresponding temporal logic (CTL) queries in the input language of the NuSMV model checker. Currently CRISNER supports dominance, consistency, preference equivalence and subsumption testing for the above languages. The obtained results from the model checker, including proofs of dominance, inconsistency or non-subsumption, are translated by CRISNER back in terms of the vocabulary of the input preference specification and saved in XML format. CRISNER’s architecture supports extension to other preference queries and preference languages such as CI-nets whose semantics are in terms of the induced preference graph.

Work in progress includes adding support for dominance and consistency testing with CI-nets, which is useful in multi-agent fair division problems. In the future, we plan to add support for computing weak order and total order extensions to consistent and inconsistent preference specifications, which is useful in applications like stable stable matching and recommender systems. We also plan to add support for reasoning with group preferences, i.e., reasoning about the preferences of multiple agents.

References

1. Abed, S., Mokhtari, Y., Ait-Mohamed, O., Tahar, S.: Numdg: A new tool for multiway decision graphs construction. *Journal of Computer Science and Technology* 26(1), 139–152 (2011)
2. Boutilier, C., Brafman, R.I., Domshlak, C., Hoos, H.H., Poole, D.: Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J.Art.Intel.Res.* 21, 135–191 (2004)
3. Boutilier, C., Brafman, R.I., Hoos, H.H., Poole, D.: Reasoning with conditional ceteris paribus preference statements. In: *UAI*. pp. 71–80 (1999)
4. Bouveret, S., Endriss, U., Lang, J.: Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In: *IJCAI*. pp. 67–72 (2009)
5. Brafman, R.I., Domshlak, C., Shimony, S.E.: On graphical modeling of preference and importance. *J. Artif. Intell. Res. (JAIR)* 25, 389424 (2006)
6. Cimatti, A., Clarke, E., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV Version 2: An OpenSource Tool for Symbolic Model Checking. In: *Computer-Aided Verification*. pp. 359–364. Springer, Copenhagen, Denmark (July 2002)
7. Clarke, E.M., Emerson, E.A., Sistla, A.P.: Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM TOPLAS* 8(2), 244–263 (1986)
8. Clarke, E., Grumberg, O., Peled, D.: *Model Checking*. MIT Press (January 2000)
9. Goldsmith, J., Lang, J., Truszczyński, M., Wilson, N.: The computational complexity of dominance and consistency in cp-nets. *J.Art.Intel.Res.* 33, 403–432 (2008)
10. Hansson, S.O.: What is ceteris paribus preference? *Journal of Philosophical Logic* 25(3), pp. 307–332 (1996), <http://www.jstor.org/stable/30226574>
11. Kripke, S.A.: Semantical considerations on modal logic. *Acta Philosophica Fennica* 16(1963), 83–94 (1963)
12. Oster, Z.J., Santhanam, G.R., Basu, S., Honavar, V.: Model checking of qualitative sensitivity preferences to minimize credential disclosure. In: *Proc. of FACS*. pp. 205–223. Springer LNCS (2012)
13. Queille, J.P., Sifakis, J.: Specification and verification of concurrent systems in CESAR. In: *International Symposium on Programming*. pp. 337 – 351. Springer Verlag (1982)
14. Santhanam, G.R.: Crisner a qualitative preference reasoner. <http://www.ece.iastate.edu/~gsanthan/crisner.html> (2015), [Online; accessed 06-February-2015]
15. Santhanam, G.R., Basu, S., Honavar, V.: Dominance testing via model checking. In: *AAAI*. pp. 357–362. AAAI Press (2010)
16. Santhanam, G.R., Basu, S., Honavar, V.: Representing and reasoning with qualitative preferences for compositional systems. *J. Artif. Int. Res.* 42(1), 211–274 (Sep 2011), <http://dl.acm.org/citation.cfm?id=2208436.2208443>
17. Santhanam, G.R., Basu, S., Honavar, V.: Verifying preferential equivalence and subsumption via model checking. In: *Algorithmic Decision Theory*, pp. 324–335. Springer (2013)
18. Wilson, N.: Extending cp-nets with stronger conditional preference statements. In: *AAAI*. pp. 735–741 (2004)
19. Wilson, N.: Computational techniques for a simple theory of conditional preferences. *Artificial Intelligence* 175, 1053 – 1091 (2011)